

RD-A151 615

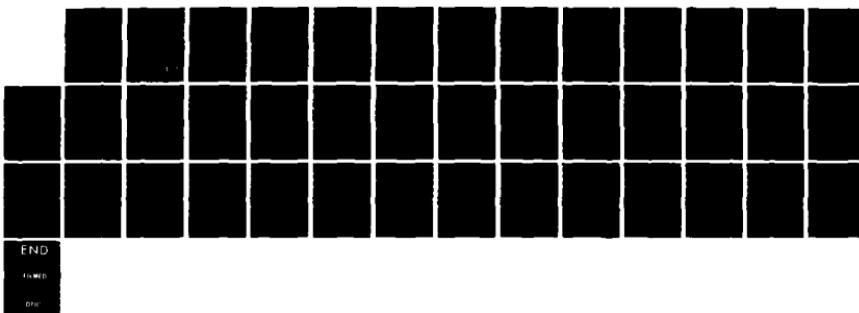
SISAL OPTIMIZATIONS(U) AIR FORCE INST OF TECH
WRIGHT-PATTERSON AFB OH J D ACRES 1984
AFIT/CI/NR-85-13T

1/1

UNCLASSIFIED

F/G 9/2

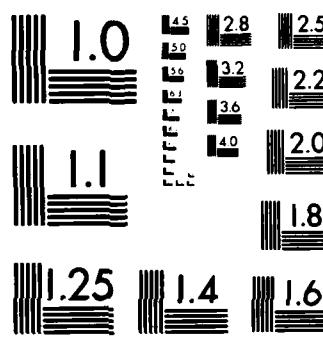
NL



END

FINISHED

DONE



MICROCOPY RESOLUTION TEST CHART
NATIONAL BUREAU OF STANDARDS-1963-A

DMC FILE COPY

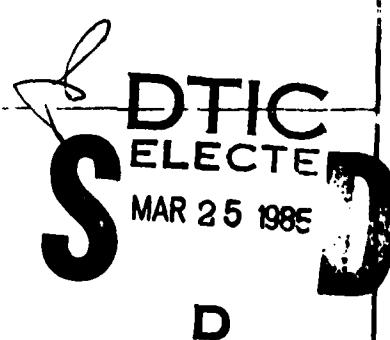
AD-A151 615

UNCLASS

(1)

SECURITY CLASSIFICATION OF THIS PAGE (When Data Entered)

REPORT DOCUMENTATION PAGE		READ INSTRUCTIONS BEFORE COMPLETING FORM
1. REPORT NUMBER AFIT/C1/NR 85-13T	2. FUNDING AGENCY ORIGIN NO.	3. EQUIPMENT'S CATALOG NUMBER
4. TITLE (and Subtitle) Sisal Optimizations		5. TYPE OF REPORT & PERIOD COVERED THESIS/DV9999XYYN
7. AUTHOR(s) Jody Dejonghe Acres		6. PERFORMING ORG. REPORT NUMBER
9. PERFORMING ORGANIZATION NAME AND ADDRESS AFIT STUDENT AT: Colorado State Univ		10. PROGRAM ELEMENT, PROJECT, TASK AREA & WORK UNIT NUMBERS
CONTROLLING OFFICE NAME AND ADDRESS AFIT/NR FB OH 45433		12. REPORT DATE 1984
MONITORING AGENCY NAME & ADDRESS (if different from Controlling Office)		13. NUMBER OF PAGES 36
		15. SECURITY CLASS. (of this report) UNCLASS
		15a. DECLASSIFICATION/DOWNGRADING SCHEDULE
16. DISTRIBUTION STATEMENT (of this Report) APPROVED FOR PUBLIC RELEASE; DISTRIBUTION UNLIMITED		
17. DISTRIBUTION STATEMENT (of the abstract entered in Block 20, if different from Report)		
18. SUPPLEMENTARY NOTES APPROVED FOR PUBLIC RELEASE: IAW AFR 190-1 LYNN E. WOLAYER 28 Feb 85 Dean for Research and Professional Development AFIT, Wright-Patterson AFB OH		
19. KEY WORDS (Continue on reverse side if necessary and identify by block number)		
20. ABSTRACT (Continue on reverse side if necessary and identify by block number)		
ATTACHED		



13

ABSTRACT

The goal of this master's project was to provide several optimizations for the SISAL compiler being developed at Colorado State University. SISAL is a data flow language intended for use on a variety of multiprocessor architectures. Since SISAL is compiled into the intermediate form IFI, which is a common intermediate form for data flow languages, this project concentrated on optimizations that are unique to the characteristics of SISAL, rather than on more traditional optimizations. In particular, the optimizations developed concentrated on making array operations more efficient by eliminating operations where possible and by performing "live analysis" which would tell if data values will ever be needed again in a program.

Accession For	
NTIS GRA&I	<input checked="" type="checkbox"/>
DTIC TAB	<input type="checkbox"/>
Unannounced	<input type="checkbox"/>
Justification	
By _____	
Distribution/ _____	
Availability Codes	
Dist	Avail and/or Special
A-1	



AFIT RESEARCH ASSESSMENT

The purpose of this questionnaire is to ascertain the value and/or contribution of research accomplished by students or faculty of the Air Force Institute of Technology (AU). It would be greatly appreciated if you would complete the following questionnaire and return it to:

AFIT/NR
Wright-Patterson AFB OH 45433

RESEARCH TITLE: Sisal Optimizations

AUTHOR: ACRES, Jody Dejonghe

RESEARCH ASSESSMENT QUESTIONS:

NAME _____ **GRADE** _____ **POSITION** _____

ORGANIZATION **LOCATION**

STATEMENT(S):

S I S A L O P T I M I Z A T I O N S

JODY DEJONGHE ACRES
CPT, USAF
FALL 1994

IN PARTIAL FULFILLMENT OF THE REQUIREMENTS FOR
MASTER OF SCIENCE
COMPUTER SCIENCE
COLORADO STATE UNIVERSITY

(36 PAGES)

1.2 OVERVIEW

The purpose of this master's project was to provide several optimizations needed for the SISAL compiler being developed at Colorado State University. In particular the optimizations considered were loop invariant removal, array optimizations and stream optimizations. The majority of the standard optimizations applied to compilers will be provided by other sources. This project's main goal, therefore, was to provide "non-standard" optimizations that are required by some of the more unique characteristics of the SISAL compiler.

2.0 BACKGROUND

SISAL (Streams and Iteration in a Single Assignment Language) is a data flow language intended for use on a variety of multiprocessor architectures. The current project at CSU is to produce a compiler for a Denelcor HEP multiprocessor. The basic characteristics of SISAL are: 1) no side effects; 2) locality of effect; 3) parallelism constrained only by data dependencies; and 4) single assignment. [Cobb, 1984] The main objective of SISAL is to provide a programming language for engineering algorithms which will make it easy to detect and exploit any implicit parallelism. The main application area being numerical computations that are straining current high performance machines. Being a "data flow" language, the statements in SISAL are not necessarily executed sequentially. Instead, they are executed whenever their input data is available. The parallelism is obtained since all instructions

whose data is available can theoretically be executed simultaneously. See [McGraw, 1983] for more information on SISAL.

The SISAL compiler produces an intermediate form referred to as "IFI". IFI is an intermediate form designed specifically for data flow languages. It describes a program as a data flow graph with nodes representing operations, edges representing data and types associated with an edge identifying the characteristics of the data being passed. Figure 1 shows an example of the graphical representation of the expression $(a * b) + z$ and the IFI description of this graph. The numbers in the boxes are port numbers that are used to identify the values. See [Brodielevski, 1984] for a complete description of IFI.

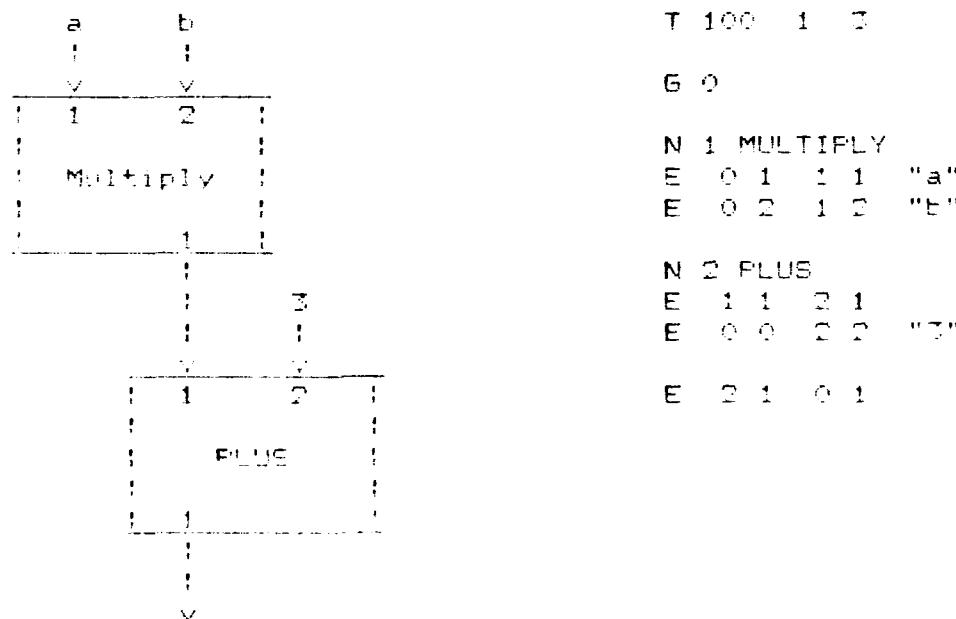


FIGURE 1

To make the IFI easier to traverse during compilation, it is encoded into a graphical format using pointers. The nodes, ports, and edges are all represented by a "C" structure containing information about each. They are linked together via pointers which allow one to traverse the graph in either a forward or backward direction. To follow the graph forward one goes from a node to a port (representing an output value), to the edge (representing a use of this value), to a node that uses the data. To follow the graph backward, one goes from a node to an edge representing an input value (physically different than the edge followed in the forward direction) to the port that this value came from, back to the node that produced the value. Figure 8 shows this graphical encoding of the previous example (A * B) + C. Going forward from a port and following the linked list of edges will give you all uses of this value. Going backward from a node and following this linked list of edges will give you all the input values for this node. The optimizations performed in this project were applied to this graphical encoding of a SISAL program.

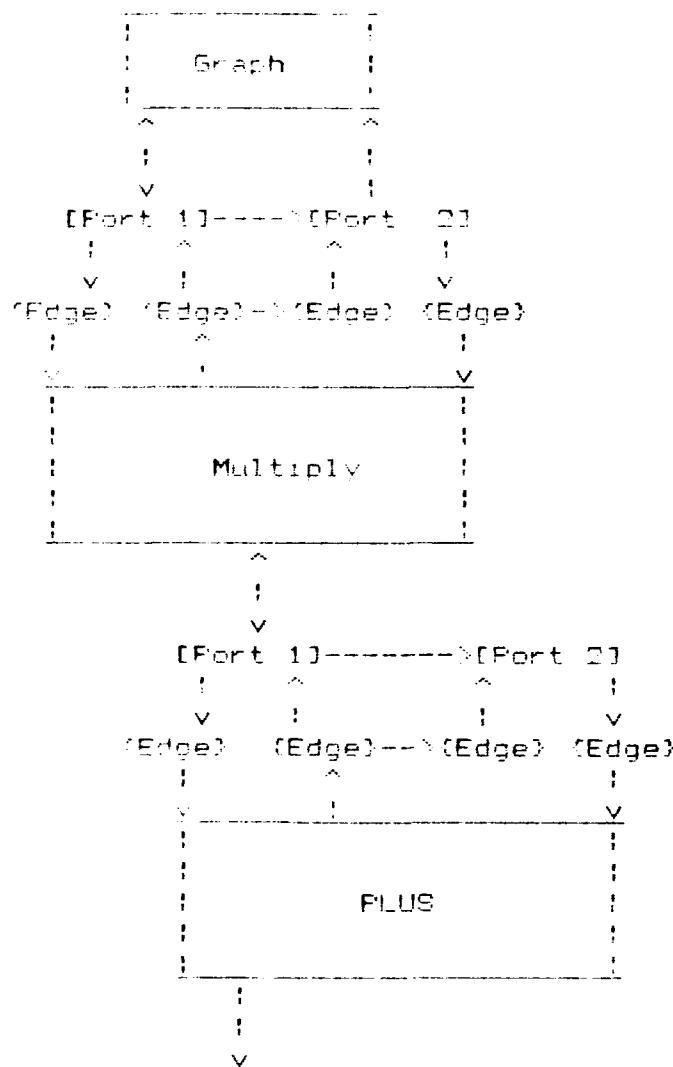


FIGURE 2

2.1. LOOP OPTIMIZATIONS

The loop optimization considered in this project was loop invariant removal. This involves finding an instruction inside of a loop whose computation stays constant for each iteration of the loop and moving it outside of the loop. This way it is only evaluated once, saving execution time. Because of the characteristics of STAL and the graphical encoding of IF1, identifying loop invariants is simpler than classical approaches to solving this problem. In particular, basic blocks and loops are already identified in the intermediate form and loop iteration information is contained in the graphical encoding. Moving the code outside of the loop, however, is more difficult.

2.2. LOOP INVARIANT REMOVAL ALGORITHM

The steps involved in removing loop invariants from IF1 are:

- 1) Traverse the graph recursively backwards and for each "while" or "forall" node found perform steps 2-6.
- 2) Traverse the nodes in the body of the loop recursively backwards marking each node as invariant if all of its input values are invariant. An input value is considered invariant if a) it is a constant; b) its source node is node 0 meaning it comes from outside the loop or c) its source node has been previously marked as invariant. If a loop node is encountered during this traversal, steps 1-5 are performed on this node before it is processed.
- 3) Once all invariant nodes have been identified, they are

moved immediately before the loop node. This may cause nodes to be renumbered.

4) The input edges to the nodes moved in Z must be modified to point to the proper source port. This will be unaffected if the source was another moved node or a constant. If the source port was from node 0, however, the corresponding sink from outside the loop must be found and the edge wired into it.

5) Any outputs from the nodes moved in Z that are still needed inside the loop must be channeled into node 0 of the loop and all nodes that use these values must be hooked up to these new ports.

6) Finally, the ports of node 0 may need to be compacted. Since the only users of a particular input value may have been nodes moved outside the loop, some ports can be deleted and the remaining ports compacted.

3.2. STATUS OF LOOP OPTIMIZATIONS

Steps 1-7 identified above had been completed when it was discovered that this process had already been done at another site. We received copies of the code and executable object and after running several tests and analyzing the code I verified that it did indeed remove loop invariants properly. It was decided to stop any further work on this effort and continue working on optimizations that were unique to the work being done at FGII. Some documentation explaining the loop optimization code we received is contained in Appdx A. Although the code compiles correctly I did have some questions about the code

itself. First, it seems to only consider an input value as transient if the source node is the graph node "node 0". In our system, however, constants have no source node. I did run some tests with constants and they were optimized correctly, so their ordering must just be different than ours. Also I didn't think the code we received was actually executable. The one thing representing levels of procedure was just incorrect. In addition, I couldn't find any place in the code that forced the inputs to nodes that weren't moved to reflect the correct location of nodes that were moved. There were, however, many routines that were not compiled with the code and this part could have been included or not of them.

ARRAY OPTIMIZATIONS

The next task was to analyze the array processing in SISAL code to determine if there were any optimizations that could be performed to make it more efficient. This involved first translating the interpretation of SISAL into IFI and then into the machine routines that actually perform the array actions to see what additional work, if any, was needed.

ARRAY OPTIMIZATIONS

The first step determines how each array operation in SISAL is mapped to an interpretation routine that actually performs the array operation.

For example, in the SISAL code shown below, with line numbers added:

 SISAL code: Interpretation code:

IFI

int	value	value
1	
1	
2	

ABUILD

array, INT

CREATE_C_BOUNDS

Parameters:

NAME	MEANING	SOURCE
nbound	number of bounds	count input edges
nvalue	number of values	count input edges
size	element size of values	type of input edge
dvflag	boolean indicating if values are arrays	type of input edge
nvadptr	new dope vector position array, of lower bounds	output param
lvadptr	array of pointers to values	input edges
rvadptr	array of pointers to input values	input edges

CreateFull - creates an array with lower and upper bounds, and
is filled with value V

CLEAR_C_BOUNDS

IFI

low	hi	value
1	1	1
1	1	1
2	1	1

AFILL

array, INT

CREATE_C_BOUNDS_FILL

Parameters:

NAME	MEANING	SOURCE
lvadptr	lower bound	current edge
rvadptr	upper bound	next edge
size	the element size	type of edge
dvflag	boolean indicating if values are arrays	type of edge
nvadptr	new dope vector position array, of lower bounds	output param
rvadptr	array of pointers to values	input edges

Select - returns a pointer to the array element at index J

SISAL - ADDJ

IF1 -

array(T)	integer	
v	v	

	AELEMENT	
v		
T		

Runtime - select

Params:

<u>NAME</u>	<u>MEANING</u>	<u>SOURCE</u>
dvptr	ptr to dope vector to start dereferencing at	from symbol table
rsize	element size	array type
ndim	number of levels in the array	array type
boundarr	ptr to subscripts to apply to array	input edges

Replacement - replaces the array element at index J with value V

SISAL - ADDVJ

IF1 -

array(T)	int	value(s)	
v	v	v	

	AREPLACE		
v			
array(T)			

Runtime - replace

Params:

<u>NAME</u>	<u>MEANING</u>	<u>SOURCE</u>
dvptr	ptr to dope vector of input array	symbol table
newdvptr	new dope vector	output param
dvflag	boolean indicating if element to replace is an array	array type
numval	number of elements to replace	input edges

<code>resize</code>	<code>element_size</code>	<code>array_type</code>
<code>ndim</code>	<code>if values are arrays,</code> <code>number of dimensions in</code> <code>these arrays</code>	<code>array_type</code>
<code>boundarr</code>	<code>array of indices to the</code> <code>element</code>	<code>input edges</code>
<code>valarr</code>	<code>array of ptrs to values</code>	<code>input edges</code>
<code>inplace</code>	<code>boolean indicating if</code> <code>input array is referenced</code> <code>again</code>	<code>?</code>

Concatenation = concatenates two or more arrays

SIGMA - ALIB

151

```

array(T) array(T)
|   |
|   v
-----
|   ADATENATE
|   |
|   v
-----
```

Planting - *cont'd*

Paradee;

<code>dvflag</code>	boolean indicating if input arrays contain dope vectors	array type
<code>numarr</code>	number of arrays being concatenated	count input edges
<code>rszsize</code>	element size	array type
<code>newdopestr</code>	new dope vector	output param
<code>dvnarr</code>	array of ptrs to dope vectors for input arrays	input edges and symbol table

Add - **High**/**Add low** - appends a single value to either end of an array.

STRA = array_add1(A, V) array_add1(A, V)

T C T -

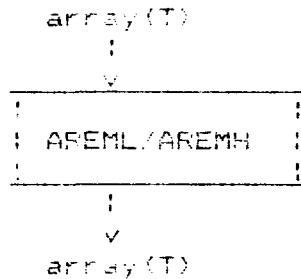
Parameters:

NAME	MEANING	SOURCE
dptr	ptr to dope vector	symbol table
newdptr	new dope vector	output param
value	data value to append	input edge
size	size of value	edge type
ifflag	boolean indicating if value is an array	edge type

Function: `high_low_element` - returns the array, A with its high index decreased by one or its low index increased by one

SISAL: `array_reml(A)` `array_remh(A)`

IF1



Runtime = `array_reml`, `array_remh`

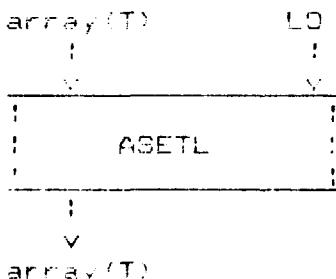
Parameters:

NAME	MEANING	SOURCE
dptr	ptr to input array	symbol table
newdptr	new dope vector	output param
size	element size of array	array type

Set low limit = adds LO + `array_lim(A)` to all elements thus shifting the origin of the array

SISAL: `array_setl(A,LO)`

IF1 -



Runtimes - array_bound

Params:

<u>NAME</u>	<u>MEANING</u>	<u>SOURCE</u>
dptr	ptr to input array	symbol table
newdptr	new dope vector	output param
lbound	lower bound for new array	input edge

Set_bounds - returns an array with range (LO,HI) with the same data as the input array where possible. If LO > array_liml(A) or HI < array_limb(A) elements will be missing

EISAL = array_adjust(A,LO,HI)

IFI =

array(T)	LO	HI
v	v	v
AADJUST		

↓
v
array(T)

Runtimes - array_adjust

Params:

<u>NAME</u>	<u>MEANING</u>	<u>SOURCE</u>
lo	lower bound	input edge
hi	upper bound	input edge
rsize	element size of array	array type
dptr	pointer to input array	symbol table
newdptr	new dope vector	output param

4.2 ANALYSIS RESULTS

There were several observations from the above analysis:

- 1) The majority of the input parameters needed by the runtime routines are directly available from either the IFI or the symbol table. The only exception to this was parameter "rreplace" used in the replace routine to determine if the input code, is evaluated again. If it isn't, the replacement can take

place without having to copy the array. Some kind of "live" node, or needs to be done to supply this value.

2) Although both SIGAL and the runtime routines allow for subscripting in referencing a multi-dimensional array, the IFI breaks them down into indexing by, only one index at a time. This is done to allow local invariant removal and common subexpression optimizations to be performed, however, it means some of the capabilities of the runtime routine will never be used. The actions taken for one call to the runtime routine versus several calls to dereference an element are probably not much different. However a routine that would combine the references into one IFI node after all optimizations have been performed could be useful.

3) A similar thing happens with concatenates as they are broken down into concatenating only two arrays at one time by the intermediate form. Both SIGAL and the runtime routine will allow multiple arrays to be concatenated at once. In this case, however, the results of several calls to the runtime routine versus just one can result in less efficiency. This happens because for each call a new logical array will be allocated to hold the results of the concatenate. One call to the runtime routine would only have to allocate one logical array. An optimization that would go through the IFI and combine any series of concatenates into one could increase the efficiency of the runtime routine.

4) There is an IFI node available called AIIsEmpty which returns a Boolean indicating if there are any elements in the array. I couldn't find any SIGAL statement that used this node.

5) There is a runtime routine which deallocates arrays. I was unable to find in, documentation that outlined when an array would be deallocated.

From these observations it was decided that routines were needed to combine a series of concatenates into one concatenate or a series of selects into one select. A routine was also decided to do a "live" analysis on an array, input to a replace node. Instead of writing a specific routine for this purpose, however, a generalized routine which given any input edge will determine if the value is ever needed again would satisfy this requirement and would also be useful in other areas of code generation.

4.2 COMBINING CONCATENATES

As mentioned above, a routine was written to combine a series of concatenate nodes in IFI into a single node when possible. Figure 7 shows an example of an IFI graph with several concatenate nodes and the graph that would result from this routine. The code for this routine is contained in appendix B and described below.

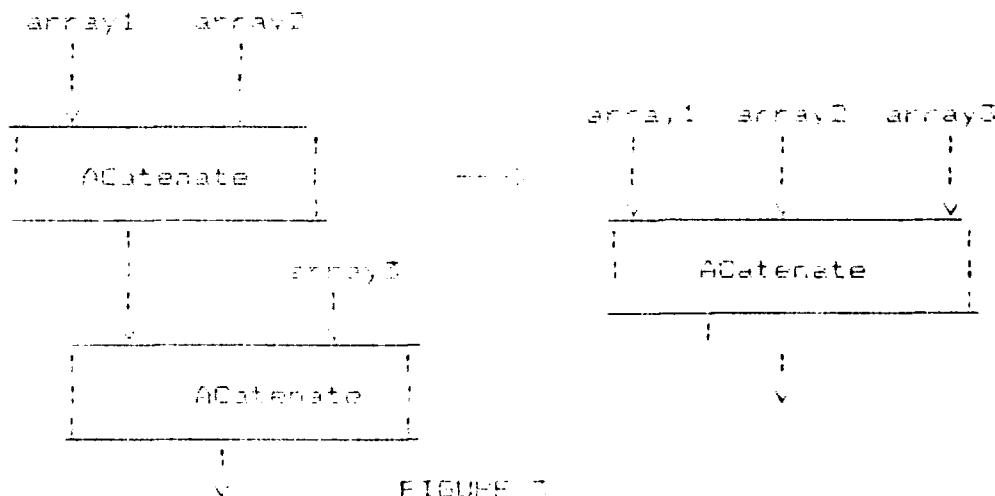


FIGURE 7

The main strategy is to start with a concatenate node from the IFI graph and then by following its inputs recursively backwards through the graph to determine if any of the inputs to this node are also concatenates. If so, the input edges of the higher concatenate are linked into the lower concatenate node instead. Before the edges are linked to the lower node, however, its inputs are also checked so that multiple levels of concatenates can ultimately be combined together.

The main routine is called "combinecats" and expects as input a pointer to a concatenate node. It first initializes a linked list which will be used to hold the input edges to the final combined concatenate node. It then calls a routine "followedges" which will cycle back through all the input edges to the node and place the inputs to all nodes that can be combined in the linked list. After it returns from this routine it sets the backptr in the original node to point to this new list of input edges and sets the next pointer of the last edge to NULL.

Routine "followedges" cycles through each of the input edges to a node. For each edge it looks at the port that produced this value and checks if there is more than one use of this value. If so, it does not combine the nodes since this intermediate value is needed somewhere else. If the value is to be used once it looks back to the source node. If it is a concatenate node, the nodes can be combined and it calls itself to start checking the inputs for this node. Once an input edge is found whose source node can not be combined the input edge is sent to the linked list of input edges for the combined node.

In addition, the corresponding "forward" edge for this input value must be made to point to the combined node instead of the original node that may have already been combined. Once all edges for the input node have been processed the routine returns. Because of the order that edges will be processed in this routine the edges placed on the linked list of edges will end up in the correct order.

Figure 4a is a detailed example of the IFI graphical encoding of two concatenates that can be combined into one. Concatenates is called with a pointer to node 5. It in turn calls followedges with node 5. Followedges looks at the first input edge. By going back up to the source port and then looking at the linked list of edges pointed to by this port it sees that there is only one use of this value. It also sees that the source node is a concatenate and so it calls itself with a pointer to node 3. The first input to node 3 is checked and although there is only one use of this value the source node is not a concatenate. This means this input edge must be placed in the linked list and the edge pointed to by node 1 port 1 must be node 4 to point to node 5 instead of node 3. Now the second input to node 3 is checked. It is similar to edge 1 and so it is also added to the linked list and its corresponding forward edge node 6 to point to node 5. This completes the processing of node 3. Next, the second input edge to node 5 is processed. It is found that the input array value is needed by another node and this node's source node can not be combined. The edge is therefore put on the linked list of edges and its forward edge node 6 to point to node 5 (which it already did).

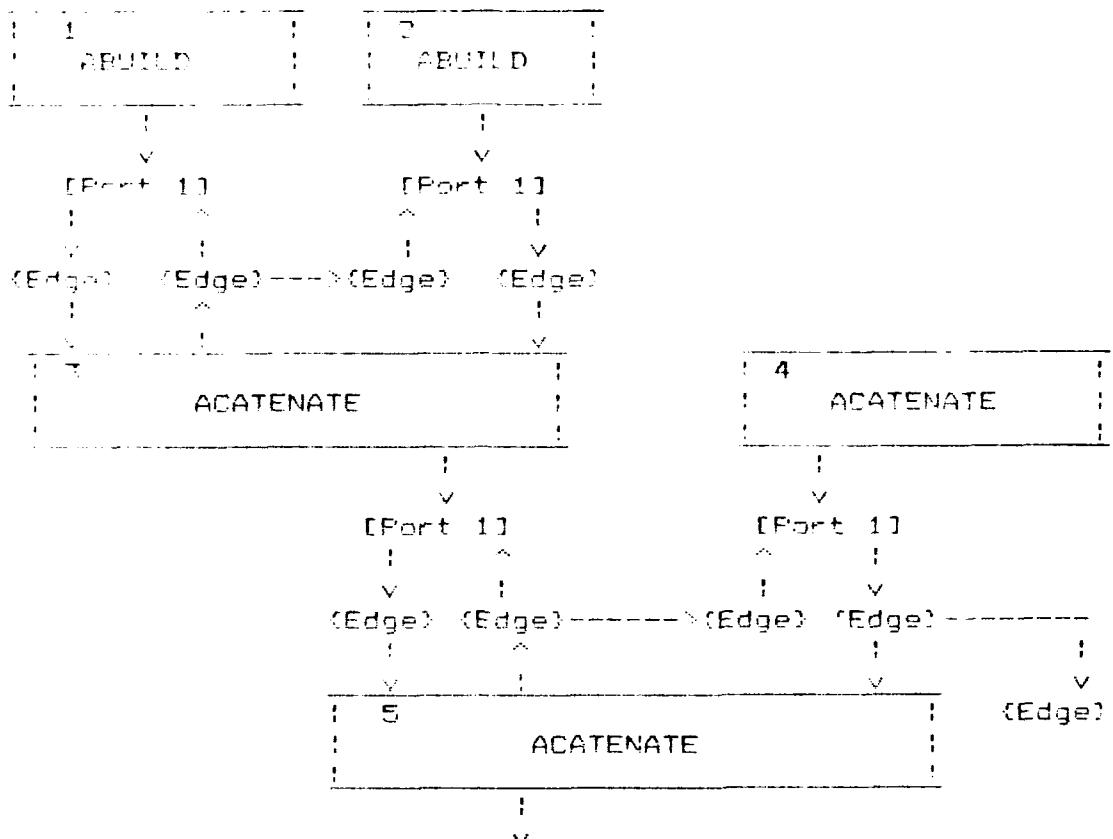


FIGURE 4a

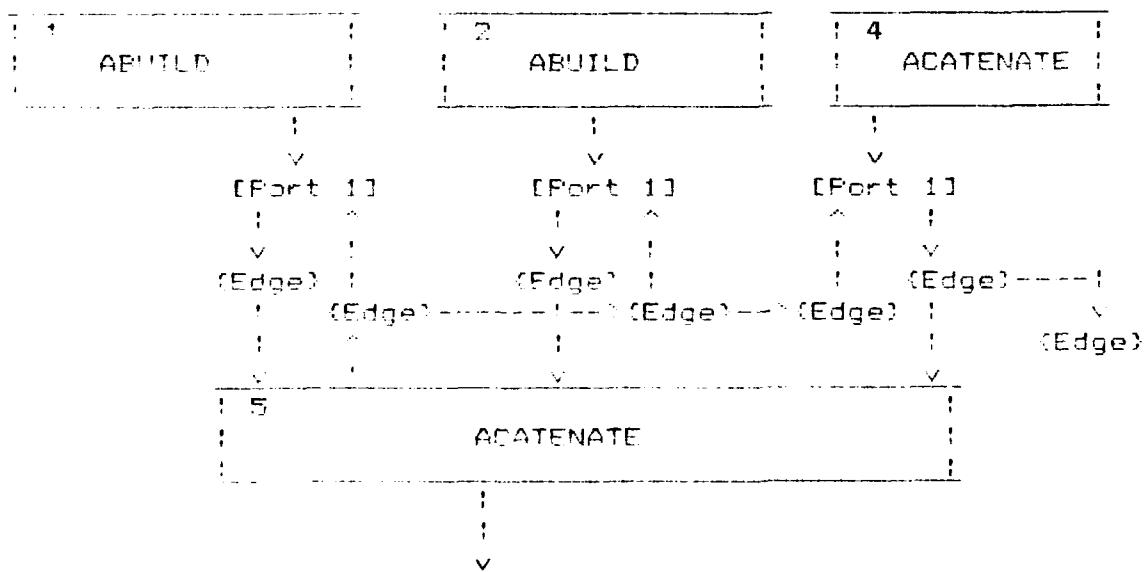


FIGURE 4b

This is the last input edge for node S, and so "followedge" returns to "combinests". The backptr for node S is set to the linked list of edges and the next pointer for the last edge set to NULL. The resulting graph is shown in figure 4b.

4.1.1 COMBINING SELECTS

The routine to combine selects is very similar to that described above for concatenator above. The goal being to construct the dereferencing of several dimensions of a multi-dimensional array into one runtime call. The code for this routine is contained in appendix C and described below.

Again, the strategy is to start with a select node from the TFI graph and then by following its inputs recursively backwards through the graph to determine if any of its inputs are "selects". If they are and there is no other use of the input value then the index for this dimension of the array is added to a "list" of indexes to be used in dereferencing an element.

The main routine is called "combineselect" and expects as input a pointer to a select or "element" node. It initializes a linked list which will hold all the indexes for the final select node. It then calls routine "followback" to cycle through all the input edges for this node and to place all the indexes that can be combined in the linked list. After it returns from this routine it fixes up the pointers to the nodes so that the nodes are the array to be dereferenced followed by all the indexes. The array to be dereferenced is the first input to the highest node in the graph that could be combined.

Routine "followback" looks at the first input to the node which is the node that is input to the select. It looks at the source node for this input and if it is also a select and there is another copy of this intermediate value, then the nodes can be combined. It then calls itself with this new node to see if the inputs to this node can also be combined. Once a node that cannot be combined is found, the routine adds the index for this node to the linked list of indexes. The corresponding "forward" edge is also made to point to the combined node. As the routine returns from all its recursive calls, all indexes will be added to the linked list.

4.5.4. THE ANALYSIS

Given a pointer to an edge, this routine will return 0 if the value is not used again or 1 if it is. The basic strategy here is to determine if there are any other uses of the value produced by the source port of this edge. Because of the parallel manner in which the IFI nodes are executed, any other use of a value could possibly be a later use of the value. The only exception would be if there was a data dependency that required the other use of this value to occur before the particular one being analyzed.

The code for live analysis is contained in Appendix D and summarized below. The main routine is called "ilive" and expects as input a pointer to the edge whose "liveness" is to be checked and a pointer to the node that this edge is input to. The edge must be the "back-edge" that one would find by following the forward edges to make out the "backward edges" one might expect by

following the edge pointer from a node. The routine first initializes *livetime* to be false (0). It then finds the point that produced the value and makes a list of all other nodes that use this value. Starting with each node on the list, it is flagged which is initially set to 1 which means that the node is alive. If there are no nodes on the list, (i.e., there are other uses), it calls "checkdepn" to see if there is a data dependency. This would examine the node using the edge in question to check whether the other nodes that use this value. If no node on the list depends on the graph's dependency, then the value is considered to be live. If no live use of this value has been found, the next step is to determine if the edge in question is embedded in a compound node and was passed in from outside the node. If so, it is impossible that there is another use of the value outside of the compound node. The graph node for the compound node and the corresponding input edge is found. The liveliness of the original edge is then determined by calling "islive" with this corresponding edge.

Protocol "checkdepn" takes as input a pointer to a node that is being checked and a pointer to a list of nodes against which it is to be checked for data dependencies. The routine cycles through each edge to the node to be checked and finds the edge that uses this value. It then cycles through the linked list of edges and sets its flag to false if a match is found. This means that there is a data dependency and the node is no longer live. If the dependent node is the graph node it immediately calls "islive" with the node to continue checking for data dependencies.

ASSAY OPTIMIZATION CONSIDERATION

TS, the last sentence is repeated above as it truly stands. The
rest of the page is blank except for the title STANLEY. This title may have
been added later by the printer who was responsible for the first printing of the
book. At the bottom of the page is a short note from the author, Stanley, giving
his address. It is written in a small hand and is very difficult to read.

There are two problems associated with the subroutine that one user should be aware of. The first is that the code is not fully combining several concatenated nodes or several pointers into effectively an illegal EIF node since it has more than two inputs. It was believed that the only use of the EIF point left after these concatenation were done would be the code generation and that this would not be a problem. Any other user who might look at this graph at this point should be aware of these pointer nodes. The second problem is in the graphical interface of EIF. Currently there is no way to jump back out of a node's code to the surrounding graph. This requires enabling the "Leave analysis" when one wants to jump out and then specifying names of the values. A pointer will have to be added to the graphical interface to allow this.

5.3 STREAM DEINITIALIZATION

The next objective of this project was to analyse the effect of variations in STGAI to determine if any significant correlations could be found with their score efficient. Likewise, this analysis will check correlations between STGAI, SEI and the other socio-economic characteristics in each community. The following table shows

exists, and if the element type is T, then stream is type T. In this case, stream is type T because they needed it. It is currently waiting for more information and thus the entire analysis could not be completed. The STREAM(IF1) relationship was analyzed and is shown in the next section.

5.1 STREAM ANALYSIS

CREATE stream - creates a stream of the desired type with no elements in it.

SIGAL <stream_type-name>T1

The stream type is created. No node is generated.

Append - appends V to the end of stream S.

SIGAL <stream_append(S,V)>

IF1

stream(T)	value
o	V
AADDH	

stream(T)

PrintTo =

Select first element - returns the first element of a stream.

SIGAL <stream_first(S)>

IF1

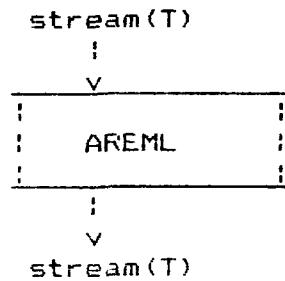
stream(T)	1
o	E
ELEMENT	

Runtime -

Select all but first element - returns a stream identical to the input stream except with the first element removed

SISAL - stream_rest(G)

IF1 -

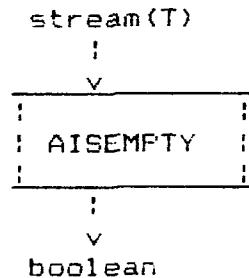


Runtime -

Test for empty - returns true if there are no elements in the stream and false otherwise

SISAL - stream_empty(G)

IF1 -

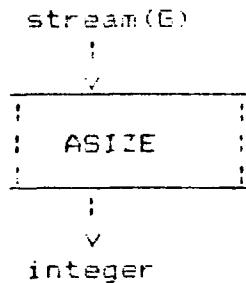


Runtime -

Number of elements - returns the number of elements in G

SISAL - stream_size(G)

IF1 -

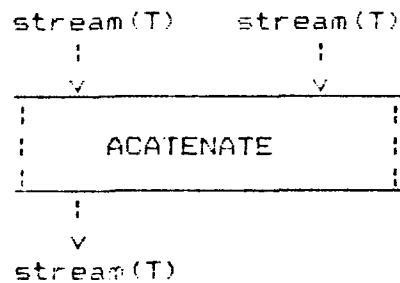


Runtime -

Concatenate - returns a stream of all elements of G followed by all elements of H

SISAL - G ++ H

IF1 -



Runtime -

5.2 STREAM ANALYSIS CONCLUSIONS

Not much can be said until the runtime routines are analyzed. It seems though that the same sort of things done for arrays may need to be done with streams. In particular, concatenates could possibly be combined if the runtime routine can handle more than two at a time. Also some kind of live analysis may be useful in determining when a stream is no longer needed.

6.2 PROJECT CONCLUSIONS

The goal of this project was to provide some optimizations for the SISAL compiler which would allow it to produce more efficient code. Since SISAL is compiled into the intermediate form IFI which is a common intermediate form for data flow languages we were looking for optimizations that were unique to the characteristics of SISAL rather than the more traditional optimizations. As we found out with loop optimization, these more traditional optimizations are being worked on at other sites. Arrays were one area of SISAL that could be a potential problem area if no optimizations were applied. In particular, replace operations could quickly eat up a lot of memory if the array had to be copied each time. The combining of concatenate and select nodes and the live analysis provided by this project will hopefully increase the efficiency of array operations in SISAL.

As for my personal gains from this project, I learned a lot about data flow machines and data flow programming languages. The concepts of data flow and parallel execution add another dimension to the way one traditionally thinks about programming. Looking at the loop optimizations also forced me to analyze those methods more thoroughly.

REFERENCES

[Cobb, 1981] Stephen F. Cobb, et. al. Arrays in SISAL. Computer Science Technical Report CS-84-04, Colorado State University, July 1984.

[McGraw, 1983] James McGraw, et. al. SISAL: Streams and Iteration in a Single-Assignment Language, Language Reference Manual. Version 1.1, July 20, 1983.

[Skadzielewski, 1984] S. Skadzielewski. IFI - An Intermediate Form for Applicative Languages. Draft 9, 1984.

APPENDIX A

IF1loop Documentation

IF1loop is a command that can be executed to perform loop optimizations on an IF1 file. This documentation describes the routines that perform this optimization. The main routine is ImproveLoop.

AddInputPorts()

This function will add input ports to a loop graph. This involves adding ports to the graph as well as its subgraphs (i.e., the test, body, and results subgraphs).

Inputs : Pointer to the node to add ports to
Number of ports to add
Output : None
Called by : MoveNodeOut
Calls : None

CompactPorts()

This routine cycles through the input ports for a loop node and deletes any that are no longer used.

Inputs : Pointer to node whose ports will be compacted
Output : None
Called by : ImproveCompoundNode
Calls : None

ImproveCompoundNode()

This function cycles through the nodes of a loop and call IsMoveable to determine if a node is invariant. If so, it call MoveNodeOut to move the node into the graph surrounding the loop. After processing all nodes in the loop, if any nodes were moved, it calls CompactPorts to renumber the inputs to the loop to reflect deletions and additions of inputs caused by the moves.

Inputs : Pointer to the compound node to be checked for invariance
Output : None
Called by : ImproveGraph
Calls : IsMoveable
MoveNodeOut
CompactPorts

ImproveGraph -

Cycles through the graph of a function looking for compound nodes. When one is found it calls itself to improve the subgraphs of this graph. It then checks whether the node is a loop construct and call ImproveCompoundNode to remove the loop invariants. After the entire graph has been processed it call RemoveGraphCSE to remove common subexpressions.

Inputs - Pointer to the first node in graph

Output - None

Called by - LoopImprover

Called - ImproveGraph

ImproveCompoundNode

RemoveGraphCSE

Immovable -

This function determines whether a node can be moved. A node can't be moved if: a) all inputs are loop constants or b) it is not associated with manipulating multiple values.

Inputs - Pointer to the node to be checked

Pointer to the surrounding graph node

Number of current inputs to the surrounding graph

Output - Boolean indicating if a node can be moved

Called by - ImproveCompoundNode

Called - None

LoopImprover -

This is the main procedure for the loop optimizations. It cycles through the functions of a program calling other routines to perform the optimizations on each function.

Inputs - Pointer to a function in the program

Outputs - Indirectly an improved IFI graph

Called by - None

Called - ImproveGraph

RenumberGraph

MovingNodeOut -

This function moves a node from one location to another by first changing the input edges to the node to reflect the environment of the outer graph. Next, AddInputPorts is called to create new input ports into the loop to hold the outputs of these moved nodes and their outputs are wired into these ports.

Inputs - Pointer to the node to be moved

Pointer to subgraph in compound node that the node to

be moved resides in
Points to compound node that the node is to be moved
out of
Output - Current number of input ports to the compound node
Called by - ImproveCompoundNode
Calls - AddInputPorts

APPENDIX B
CODE TO COMBINE CONCATENATE NODES

```
#include "list.hf.encode/syndef.h"
struct edgetype *eliststart, *elistend;
/* Used to keep a linked list of input edges for the */
/* combined node */
struct IPnode *mesternode;
/* Put *s to the combined node */

*****COMBINECATS*****
/* This routine will combine a series of concatenate nodes */
/* by following the input edges of a concatenate node and */
/* adding linking the input edges of all nodes that can be */
/* combined into one node. */
*****COMBINECATS*****

followedges(nodeptr)
struct IPnode *nodeptr; /* Points to node to start */
/* combining at */

{
    mesternode = nodeptr; /* Remember this starting node */
    eliststart = elistend = NULL; /* Initialize list of edges */
    edges = 1; /* Used to renumber edges */
    full_wedge(nodeptr); /* Call followedge to do the */
    /* actual combining */
    nodeptr->backptr = eliststart; /* Make node point to this */
    /* new list of edges */
    elistend->nextedge = NULL; /* Fix last edge on list */
}

*****FOLLOWEDGES*****
/* This routine will cycle recursively backwards through */
/* a graph determining if nodes can be combined. When it */
/* reaches the last node that can be combined on a given */
/* path it puts its input edges on a list of edges for the */
/* final combined node. */
*****FOLLOWEDGES*****


followedges(nodeptr)
struct IPnode *nodeptr; /* points to the node whose edges */
/* are to be followed */

{
    struct edgetype *sourcededge, *destedge;
    struct porttype *sourceport;

    sourcededge = nodeptr->backptr; /* get first input edge */
    /* now cycle through all its input edges checking if the */

```

```

/* input: source node can be combined */

while (sourcedge != NULL) {
    sourceport = sourcedge->ptr.up;
    /* if this value is not needed elsewhere and the source */
    /* node is a concatenate it can be combined */
    if ((sourceport->usage->nextedge == NULL) &&
        (sourceport->source->tynode == catenate)) {
        sourceport->source->beenhere = 1;
        followedges(sourceport->source);
    }
    else { /* the source node can not be combined and */
        /* this edge is put on the edge list */
        if (eliststart == NULL)
            eliststart = sourcedge;
        else
            elistend->nextedge = sourcedge;
        elistend = sourcedge;
        /* the corresponding "forward edge" must be made to */
        /* point to the final combined node */
        destedge = sourcedge->ptr.up->usage;
        while (destedge->ptr.bn != nodeptr)
            destedge = destedge->nextedge;
        destedge->ptr.bn = masternode;
        destedge->port = sourcedge->port = edgeno++;
    }
    sourcedge = sourcedge->nextedge;
}

```

APPENDIX C
CODE TO COMBINE SELECT NODES

```
#include "sicalrifencoded/symdef.h"
struct edgeType *eliststart, *elistend;
/* Used to keep a linked list of input edges for the */
/* combined node */
struct IFnode *masternode;
/* Points to the combined node */
struct IFnode *topnode;
/* Points to the highest node in the graph */
/* to be combined */

*****COMBINESEL*****
** This routine will combine a series of select or
** "element" nodes by following the array input to the
** node and linking the indexes of all nodes that can be
** combined into one node.
*****
```

int combineSel(nodeptr)
 struct IFnode *nodeptr; /* Points to node to start */
 /* combining at */

 masternode = nodeptr; /* Remember this starting node */
 topnode = nodeptr; /* initialize the highest node */
 /* to be the current node */
 eliststart = elistend = NULL; /* Initialize list of indexes*/
 edgeno = 0; /* Used to renumber edges */
 followback(nodeptr); /* Call followback to do the */
 /* actual combining */
 nodeptr->backptr = topnode->backptr; /* make node point to */
 /* input array */
 nodeptr->backptr->edge = eliststart; /* link in list of */
 /* indexes */
 elistend->nextedge = NULL; /* fix last edge on list */
 /* now fix up first input edge */
 edgeptr = topnode->backptr->ptr.up+usage;
 while (edgeptr->ptr.dn != topnode)
 edgeptr = edgeptr->nextedge;
 edgeptr->ptr.bn = nodeptr;
 edgeptr->port = 1;

```

***** FOLLOWBACK *****
/*
 * This routine will cycle recursively backwards through
 * a graph determining if nodes can be combined. When it
 * reaches the last node that can be combined on a given
 * path it puts its index on a list of indexes for the
 * final combined node.
*/
***** FOLLOWBACK *****

followback(nodeptr)
struct TNode *nodeptr; /* points to the node whose edges */
/* are to be followed */
{
    struct edgetype *sourcededge, *destedge;
    struct porttype *sourceport;
    /* now check if the input's source node can be combined */
    sourceport = nodeptr->backptr->ptr.up;
    /* if this value is not needed elsewhere and the source */
    /* node is an element it can be combined */
    if ((sourceport->usage->node == NULL) &&
        (sourceport->usage->typenode == anna,element)) {
        to_node = sourceport->source;
        if (node->beenhired == 1)
            followback(tonode);
    }
    /* now edge = nodeptr->backptr->nextedge;
     * put the index on the linked list */
    if (eliststart == NULL)
        eliststart = sourcededge;
    else
        eliststart->nextedge = sourcededge;
    elistend = sourcededge;
    /* The corresponding "forward edge" must be made to */
    /* it point to the final combined node */
    destedge = sourcededge->ptr.up->usage;
    while (destedge->ptr.id != nodeptr)
        destedge = destedge->nextedge;
    destedge->ptr.id = masternode;
    destedge->port = sourcededge->port + edgenext++;
}
sourcededge = sourcededge->nextedge;

```

APPENDIX D

CODE TO TEST IF AN EDGE IS LIVE

```

#include "sisal/ifencode/symdef.h"
struct nodelist {
    struct IFnode *nodeptr;
    int depends;
    struct nodelist *nextnode;
} /* This structure will be used to hold a list nodes that */
/* also use the value being tested */

***** ISLIVE *****
/* This routine checks if a given data value (represented by */
/* L) an edge will ever be used again. */
***** ISLIVE *****

live(innode,inedge)
    struct IFnode *innode      /* points to the node that the */
                                /* edge in question is input to*/
    struct edgetype *inedge;   /* points to the "backEdge" */
                                /* whose liveness is to be */
                                /* tested */

    struct nodelist *nliststart, *nptr; /* contains a linked */
                                /* list of other nodes that use */
                                /* this data value */
    struct edgetype *edgeptr;
    int live, found;

    live = 0;           /* initialize live to false */
    /* build list of all other uses */

    nliststart = NULL;
    edgeptr = inedge->ptr.up_usage;
    while (edgeptr != NULL) {
        if ((edgeptr->ptr.dn != innode) ||
            (edgeptr->port != inedge->port)) {
            nptr = (struct nodelist *) (malloc(sizeof
                (struct nodelist)));
            nptr->nodeptr = edgeptr->ptr.dn;
            nptr->depends = 1; /* initialize it to be live */
            nptr->nextnode = nliststart;
            nliststart = nptr;
        }
        edgeptr = edgeptr->nextedge;
    }
}

```

* If the node is not found, then the node is added to the list of nodes, *
* and the search continues.

11. If the node is found, then the node is added to the list of nodes, *
* and the search continues.

* If the node is not found, then the node is added to the list of nodes, *
* and the search continues.

12. If the node is found, then the node is added to the list of nodes, *
* and the search continues.

13.

* If the node is found, then the node is added to the list of nodes, *
* and the search continues. If not, then the value was packed into the
* structure from the node. Find the corresponding edge and add
* it to the list of edges. *

14. If the edge is not found, then the node is added to the list of edges, *
* and the search continues. If so, then the edge is added to the list of edges, *
* and the search continues.

* If the node is found, then the node is added to the list of nodes, *
* and the search continues. If not, then the value was packed into the
* structure from the node. Find the corresponding edge and add
* it to the list of edges. *

15. If the node is found, then the node is added to the list of nodes, *
* and the search continues.

while (found) {

 if (edgeptr->parent == nodeptr->ptr) {
 found = 1;

 else
 edgeptr = edgeptr->nextedge;

16. If not, then add the edge to the list of edges.

17. If not, then add the node to the list of nodes.

DEFINITION

* This routine will check if a node has any sets
* of dependencies that would not be detected by the
* dependency checker. This is done by building a linked list
* of all output edges to nodes. If there is an edge in the output list
* to the node, then the node is added to the list of nodes.

18. If the node is found, then add it to the list of nodes.

19. If the node is not found, then the node is added to the list of nodes.
* If the node is found, then the node is added to the list of nodes.

20. If the node is not found, then the node is added to the list of nodes.
* If the node is found, then the node is added to the list of nodes.

21. If the node is found, then the node is added to the list of nodes.

END

FILMED

4-85

DTIC